DOTSE Report 172
NR 1350
ISSN 1174-3387

FIREPOWER CONCENTRATION IN CELLULAR AUTOMATA MODELS — AN ALTERNATIVE TO THE LANCHESTER APPROACH

M K Lauren

June 2000

DOTSE Report 172
NR 1350
ISSN 1174-3387

FIREPOWER CONCENTRATION IN CELLULAR AUTOMATA MODELS

— AN ALTERNATIVE TO THE LANCHESTER APPROACH

M K Lauren

June 2000

Abstract:

The ISAAC cellular automaton combat model appears to behave in one of two qualitatively different ways: either the automata fight as a massed force, or they form dispersed fractal-like distributions. The latter has properties not described by the Lanchester equation, owing to the ability of the automata to concentrate their firepower. This behaviour has a non-intuitive effect on the relationship between kill probability of the automata and attrition rate, providing insight into the possible ways a dispersed force may fight, and leading to a markedly different formulation of combat attrition from the more conventional description suggested by Lanchester. In particular, certain ensembles of runs are characterised by one force doing substantially better than might be expected from a Lanchester model. This improvement is not due to chance, but rather due to the fact that one force has been able to concentrate its forces in an opportune way. The results suggest that this characteristic of dispersed forces may allow poorly armed forces to occasionally overcome superiorly armed forces, particularly if that force does nothing to prevent itself from being concentrated upon.

Defence Operational Technology Support Establishment
Auckland, New Zealand

EXECUTIVE SUMMARY

Background

In 1914 Lanchester laid the foundation for many of the mathematical analysis methods used to describe warfare this century. He did this by introducing a pair of equations dependent on the numerical strength of the opposing forces, and the killing rate of individuals in each force. The solutions to these equations are used as the basis for many modern combat models and studies. Furthermore, most analytical combat models, which model battles by representing individual participants and the physics of their equipment, strongly resemble probabilistic versions of Lanchester's model.

The main weakness with this approach is that it treats the combatants as cannon fodder, the objective being simply to grind the opponent down by brute force. Noting the coincidence of the Lanchester equations and the bloody and futile four years of European land war from 1914-1918, it is tempting to say that their development was a retrograde step from which military analysis is yet to recover.

Recently, New Zealand's Defence Operational Technology Establishment has been working with the United States Marine Corp as part of a collaborative effort called Project Albert. This project is dedicated to the application of complex adaptive system models and associated analytic techniques to decisions made by the Marine Corp and the New Zealand Army.

In applying the methods of cellular automata to combat modelling, it is hoped to consider what happens to the possible combat outcomes when the participants are allowed to react to each other. The current model in use here is called ISAAC (Irreducible Semi-Autonomous Adaptive Combat), and was developed for the Marine Corp Combat Development Command, by the Center for Naval Analyses. This report explores the nature of the ISAAC model, particularly in regard to the way it models dispersed forces, something that the Lanchester approach is unable to do adequately.

Sponsor

New Zealand Army.

Aim

To demonstrate the ability of cellular automata models to describe the modern dispersed battlefield.

Results

Some significant differences between cellular automaton and Lanchester models are identified, particularly in regard to modelling manoeuvre warfare.

It was found that by introducing automata "personalities" to the entities being modelled, a dispersed force would adopt a fractal distribution, reflecting the degree to which firepower was concentrated. Such a distribution is neither uniformly dispersed, nor tightly concentrated in space.

The Red automata in the cases examined were able to concentrate their firepower to improve the rate of attrition of the Blue force in circumstances where they lacked the firepower to destroy Blue on contact. At the same time, if Red needed to concentrate its forces for long periods of time to achieve this, Blue would have the opportunity to exploit the gaps created in the battlefield.

Selecting just those runs for which Red was able to cause some predetermined level of Blue casualties, the attrition rate for this ensemble obeyed a power-law dependence on kill probability. The power law exponent was apparently related to the fractal dimension of the distribution of forces. The value of the power-law exponent was typically around 1/3. By contrast, a stochastic Lanchester model produced a power-law exponent of nearly 1. Thus this ensemble does better as the value of the kill probability of its individual units is decreased than would be expected from a Lanchester model, because the units have concentrated their firepower in a beneficial way. What is so interesting about this is that the superior performance of this ensemble is not entirely due to luck (as it would be for a stochastic Lanchester model), but due to the way the automata have distributed themselves for each run in this ensemble.

This suggests a poorly armed Red force can cause considerable damage to a superiorly armed Blue force if it is able to appropriately concentrate its forces. This is a particularly important consideration for peacekeeping activities in the Third World, where a poorly armed or trained militia, which may not seem like much of a threat, can cause a disaster for a well-equipped modern force. On the other hand, if the Blue force is able to react to such a concentration of force, it will often cause firepower "holes" in the Red distribution which it can exploit. Doing so will allow it to perform much better in a typical run. However, the distribution of outcomes becomes more complicated. Giving both sides the ability to react to each other's attempts to concentrate firepower leads to a much wider range of outcomes for the model, and may be somewhat bimodal, with Blue either being caught by a Red firepower trap, or being able to create firepower gaps which it can exploit. In short, treating one side (or both) as a "lifeless mass" fails to produce the full range of possible outcomes.

Contents

1. Introduction

In 1914, Lanchester[1] laid the foundation for many of the mathematical analysis methods used to describe warfare this century. He did this by introducing a pair of coupled ordinary differential equations:

$$\begin{aligned} \frac{dR}{dt} &= -k_1 B(t), \quad R(0) = R_0 \\ \frac{dB}{dt} &= -k_2 R(t), \quad B(0) = B_0 \end{aligned} \tag{1.1}$$

where R and B represent the numerical strength at time t of opposing Red and Blue forces, and k_1 and k_2 the killing rate of a Red/Blue individual. The solutions to these equations or numerical versions of these are used as the basis for many modern combat models and studies[2,3,4,5]. Furthermore, most analytical combat models, which model battles by representing individual participants and the physics of their equipment, strongly resemble quantised, stochastic versions of Lanchester's model[6].

The main weakness with this approach is that it treats the combatants as cannon fodder, the objective being simply to grind the opponent down by brute force. Noting the coincidence of equations 1.1 and the bloody and futile four years of European land war from 1914-1918, it is tempting to say that their development was a retrograde step from which military analysis is yet to recover. Indeed, knowledge of these equations failed to prevent the combined forces of France and Britain from being overrun by a German Army still recovering from years under the restrictions of the Treaty of Versailles, and were virtually useless for the Western forces in Vietnam. Why in recent times the austere mathematics of the Lanchester equation and its modern offspring, the analytical combat model, seem to have been preferred to the basic principles of war developed over the last two-and-a-half thousand years, and described as early as Sun Tzu's "The Art of War", and demonstrated by Alexander the Great, Napoleon Bonaparte, and more recently, George Patton and the Viet Cong, is something of an enigma. This appears to have been recognised by a growing number of military writers (e.g. the US Marine Corp's Lt. Gen. J. E. Rhodes[7], Lt. Col. J. F. Watson of the UK's Combined Arms Training Centre Battle Group[8], and Maj. Gen. R. H. Scales Jr, US Army War College[9], to name some examples).

Commanders such as Napoleon appear to have had a strong feel for manoeuvre and the importance of troop quality, factors which have proved vital on the battlefield, but which are neglected by the Lanchester equation. They are particularly important on the modern battlefield, where high-lethality weapons have forced troops to be deployed in a much more dispersed and fluid manner. In this report, we examine the use of cellular automaton models to describe combat, highlighting the difference between these and the Lanchesterian approach.

2. Cellular automaton model

Recently, New Zealand's Defence Operational Technology Establishment has been working with the United States Marine Corp as part of a collaborative effort called Project Albert. This project is dedicated to the application of complex adaptive system models and associated analytic techniques to decisions made by the Marine Corp and the New Zealand Army.

Our early efforts centre on the application of a cellular automaton model to combat modelling. In applying the methods of cellular automata to combat modelling, it is hoped to consider what happens to the possible combat outcomes when the participants are allowed to react to each other. Other workers have already shown that non-intuitive trends emerge from stochastic Lanchester models simply by splitting them into different arrangements of smaller engagements[10]. Automaton models extend upon this by possessing the ability to make decisions about how and when to split, and for a given run the "personalities" of the automata generally determine when the fight is over. By contrast, questions such as how the forces split and what the withdrawal level is are awkward for the Lanchester approach, because they depend on the arbitrary decision of the modeller. Consequently, automaton models are becoming an increasingly popular tool for combat modelling[11].

The model used here is called ISAAC (Irreducible Semi-Autonomous Adaptive Combat), and was developed for the Marine Corp Combat Development Combat by the Center for Naval Analyses.

The model is relatively straight-forward and only a brief description will be given. It is described in detail in Ilachinski[12], and the executable program may be downloaded from the Web address www.cna.org/ISAAC. For the simulations here, the automata have a single set of characteristics for each side, described by parameters which break down into three classes: attributes, personalities and meta-personalities. The first class describes movement rate, weapons range, weapons kill probability, and sensor range. The second class are weightings which describe an automaton's propensity to move toward/away from friendly/enemy automata, and toward/away from a goal (flag) point. The model calculates automata moves by summing the number of friendly and enemy automata within a threshold range of each square within movement range, and uses the personality weightings to determine the "penalty" associated with each move in terms of how it positions the automaton relative to the friendly/enemy forces and the goal point. Communications may also be simulated, by including the number of friendly/enemy automata visible to other friendly automata within a set communications range.

The third class modifies the above procedure. Three such parameters were used: the cluster parameter, which "turns off" an automaton's propensity to move towards friendly automata once a threshold cluster size has been reached; the advance parameter, which requires a threshold number of friendly automata to be surrounding an automaton before it will advance towards its goal; and a combat parameter, which will only allow an automata to advance towards the enemy once a threshold numerical advantage has been achieved.

The automata themselves can be in one of three states: alive, in which case they use the baseline parameters; injured, where a secondary set of parameters may be used to indicate the automata has suffered damage; and dead, in which case the automaton is removed from the battlefield.

3. Behaviour of the model and application

The principal difficulty in applying the ISAAC model to real questions of interest to the New Zealand Army and the Marine Corp is identifying the types of situations where automata models provide useful results that traditional methods do not. Qualitatively examining the behaviour of the model, the automata appear to fight in two distinct ways (see Figure 1): I) by clumping together, or II) fighting in a dispersed, fluid pattern[6].

The two cases shown in Figure 1 are I) two equal forces attempting to capture a “flag” at the centre of the battlefield, and II) a numerically inferior Blue force (20 automata) trying to break through a dispersed and numerically superior Red force (40 automata). Case I is representative of the concept of “linear” warfare (i.e. armies fighting in very rigid formations such as lines or columns), while case II is a much more fluid situation reminiscent of manoeuvre warfare. Case I is no longer considered a feasible method of warfighting, due to the high lethality of modern weapons. Having forces in such high concentrations risks making them easy targets. In manoeuvre warfare, the emphasis is on using mobility to destroy an opponent’s key points or otherwise disadvantaging him, rather than destroying him by attrition alone.

This paper deals with variations of case II. We imagine that the Blue force must manoeuvre its way through a dispersed Red force to strike some vulnerable point. As such, the focus here is on Blue success, rather than how much damage Blue causes to Red. It is imagined that if Blue achieves its goal, it will cause sufficient damage to Red’s command and supply infrastructure as to be disastrous.

At a glance, one might expect that the motion of the automata is Brownian for case II, leading to an on-average uniformly random distribution of Red forces. However, the personality rules of the simulation impose structure on the distribution as automata are attracted towards/repelled by friendly/enemy automata. This leads to groupings of automata into clusters which may “cooperate” in attacks on opposing automata.

Three variations on case II are considered: i) the Red force does not react to the Blue force; ii) Red forces within detection range of Blue will cluster around it as it attempts to break through; iii) the Blue force will react to the Red force as it surrounds it. The parameters used for each set of automata are given in the appendix.

Case iii is shown in Figure 1. The initial distributions of the Red and Blue forces are random, but confined within a finite area. As Blue enters the area containing the Red force, Red reacts by clustering around it. For this case, Blue has a weighting to move towards Red, so that its lead elements can attack and push back nearby Red forces. This quickly leads to a dispersed battlefront.

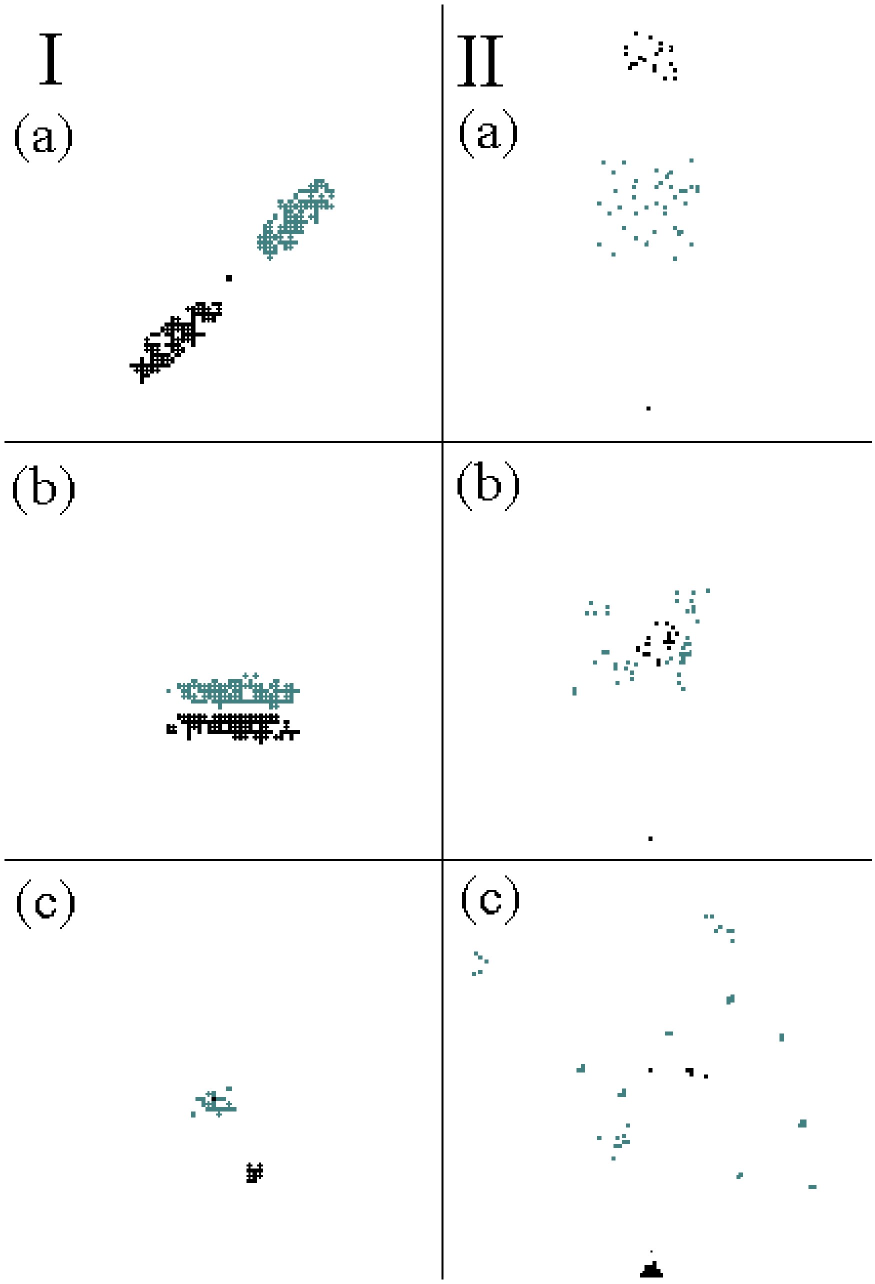

Figure 1: Evolution of the ISAAC model demonstrating the difference between the "clumpy" and "dispersed" cases.

One method for describing the degree of clustering on the battlefield is to obtain a fractal dimension for the distribution. This is done using a box-counting technique, finding the smallest square containing all the automata, and splitting this into four equal squares and counting the number of squares containing automata. Then each of these new squares is split in four, with automata-containing squares counted again, and so on. This is illustrated schematically in Figure 2. Note that this is done just for the Red Forces, since for the analysis that follows, we shall be mainly concerned with the concentration of Red firepower. If the distribution is fractal, then the number of squares containing automata should exhibit a non-integer power-law dependence on the size of the squares used.

The "box-counting" fractal dimension[13], D, defined here as:

$$D = \lim_{d \to 0} \frac{\log N}{\log\left(\frac{1}{d}\right)} \tag{3.1}$$

(where d is the width of the box, and N the number of boxes required to cover all the automata), is a power-law exponent, since:

$$N(d) = d^{-D} \tag{3.2}$$

The application of such techniques to automata models was pioneered by workers such as Bak and Chen[14,15], who showed that non-integer power-laws characteristic of fractal distributions emerged from sets of automaton rules such as the Game of Life[16,17]. Obtaining the fractal dimension allows a measurement of the degree to which one side concentrates its forces, while still remaining somewhat dispersed. A force which is spread uniformly across the battlefield will have D = 2 (using the definition in 3.2). Conversely, an extremely concentrated force will tend towards D = 0, which represents the force occupying a single point. A fractal distribution represents an "in-between" state where the force is neither tightly bound nor uniformly dispersed. Furthermore, fractals can generally be viewed as patterns which have structure on all scales. From the point of view of the distribution of the forces, the existence of a non-integer D implies that the automata form clusters which can be viewed as a collection of sub-clusters, which themselves are collections of smaller clusters, and so on, rather than simply being a single clump of automata.

In order to accurately obtain a fractal dimension representative of the distributions which evolve from a particular set of parameters, a large ensemble of distributions must be used, and an average D found. Figure 3 shows how the mean number of boxes containing automata varies with the width of the box for case iii. The first two points from the right hand side indicate the cases where the automata occupy all the boxes, and are not part of the fractal scaling range.

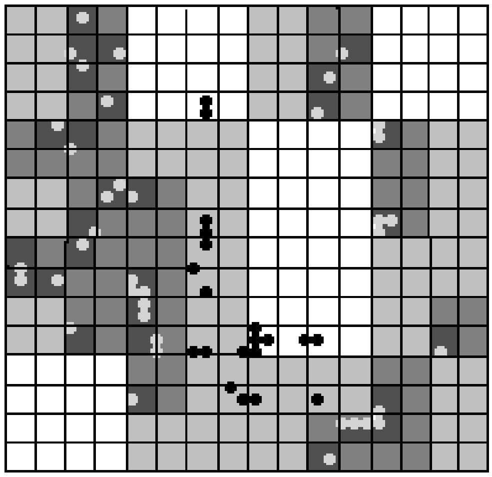

Figure 2: Schematic demonstration of a box-counting technique to determine the fractal dimension of the Red distribution. The number of boxes required to cover all the Red automata are counted as the size of the boxes is reduced.

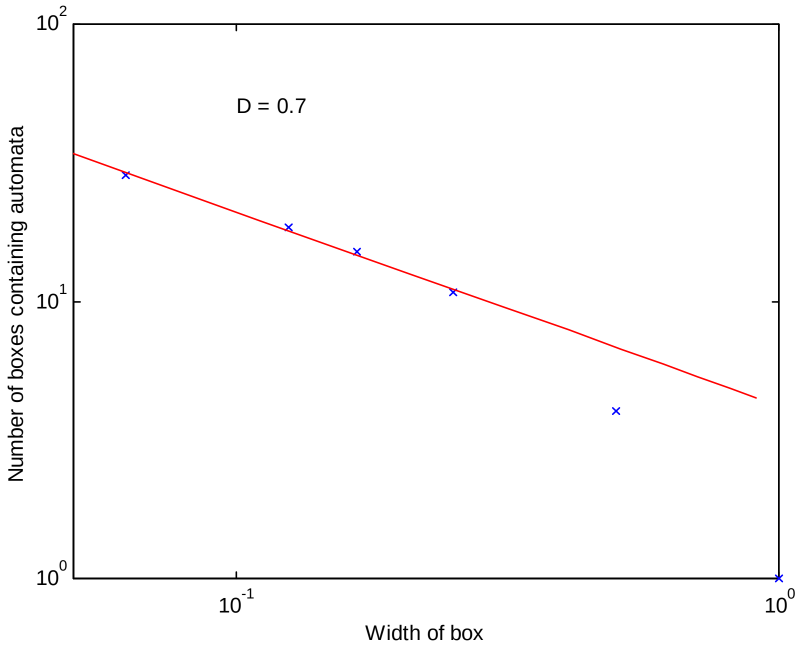

Figure 3: The non-integer power-law dependence of the number of boxes required to cover all automata on the width of the box.

Likewise, as the width of the boxes approaches the situation where all boxes contain no more than one automaton, the power law breaks down. This leaves a short scaling range in between these two extremes. D was estimated as 0.70 for case iii.

The attrition rate of the Blue force for a given run clearly depends on the degree to which Red is able to concentrate its firepower. Ideally, we would like to find a relation between D and the attrition rate. Consider the following argument: In a fixed time interval, a Blue automaton's "sensors" on average map out an area (box) proportional to the speed it is travelling. The density of Red automata if they are encountered is related to the number of boxes of this size that actually contain automata. But as seen in Figure 3, this has a non-integer power law dependence on the size of the box, and hence on the time interval (i.e. $l \propto t$). In such a case, the attrition rate function depends on:

$$\frac{\Delta B}{\Delta t} = f(k_{red}, \Delta t^{-D} n_{box}) \tag{3.3}$$

where the variable R in equation 1.1 has been replaced by $\Delta t^{-D} n_{box}$, and n_{box} is the average number of Red automata contained in just those boxes that contain automata, where the size of the box is the area an automaton typically maps out in a period Δt. If there exists an ensemble of runs for which Blue meets Red in concentrations proportional to $\Delta t^{-D} n_{box}$, it may be hypothesised that for this ensemble $\Delta B / \Delta t$ is a function of time to a non-integer power, and kill probability. If these are the only two variables on which $\Delta B / \Delta t$ depends, then since k has units of $(\text{time})^{-1}$, dimensional arguments require that for this ensemble:

$$\left\langle \frac{\Delta B}{\Delta t} \right\rangle \sim k^{q(D)} \, \Delta t^{r(D)} \tag{3.4}$$

where r is a non-integer power related to D, the angled brackets denote an ensemble average, and the ensemble is conditional on Blue encountering Red in clusters of size n_{box}.

4. **Role of force concentrations in determining outcome**

Figure 5 shows the probability densities for the number of Blue casualties suffered by the 500^{th} time step of the model for variations i, ii and iii of case II. The k values were set so that the mean casualty level was about 15 per cent, and each distribution represents 2000 runs. The important point to realise is that by this time step (and indeed, well before), the battle has run its course, so that further casualties are unlikely (i.e. Red and Blue are no longer interacting). Thus, casualties do not necessarily reach 100 per cent (or some other arbitrary level), because the Blue force often breaks through to its goal point, or otherwise avoids encountering concentrations of Red, before sustaining heavy casualties. Alternately, Blue's attack may be broken up, with the remnants avoiding Red but failing to reach their goal.

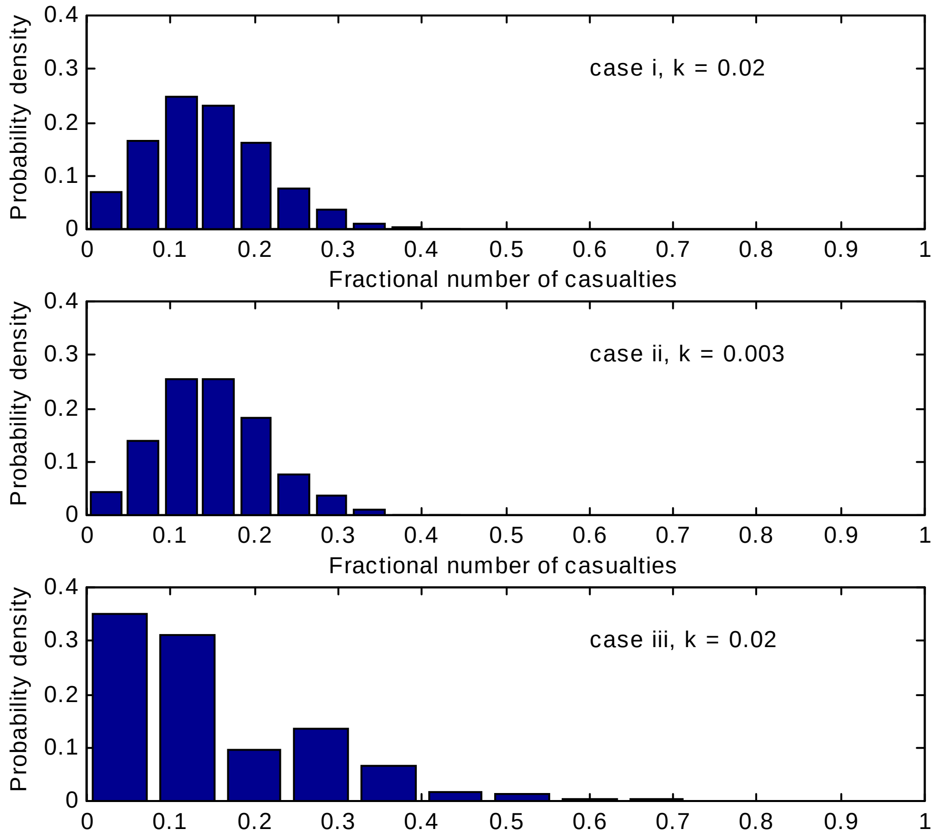

Figure 4: Probability densities for the three variations on case I. Each has a mean casualty level of 0.15 per cent.

This behaviour results from the personalities given to the automata. This contrasts with stochastic Lanchester models, where a spread of outcomes is expected after a given period of time, but the level of casualties at that point represents a “snapshot” of casualty levels on the way to some arbitrary withdrawal level.

Figure 4 thus represents the probability of ending a run in a particular state, defined by final casualty levels. Clearly, Blue only reaches a given level of casualties for a certain ensemble of runs. These runs represent the occasions where Red was able to concentrate its firepower sufficiently to inflict this level. We may explore the nature of this ensemble by measuring attrition rate in terms of time taken to reach that level of casualties, excluding runs which do not.

Figure 5 plots attrition rate as a function of k, calculated in this way, for the ensemble of runs with greater than 25 per cent Blue casualties. Lowest limits for k were set by a requirement for more than a few per cent of the ensemble to reach the 25% Blue casualty level. Cases i to iii are as described above, case iv is the linear case (case I shown in

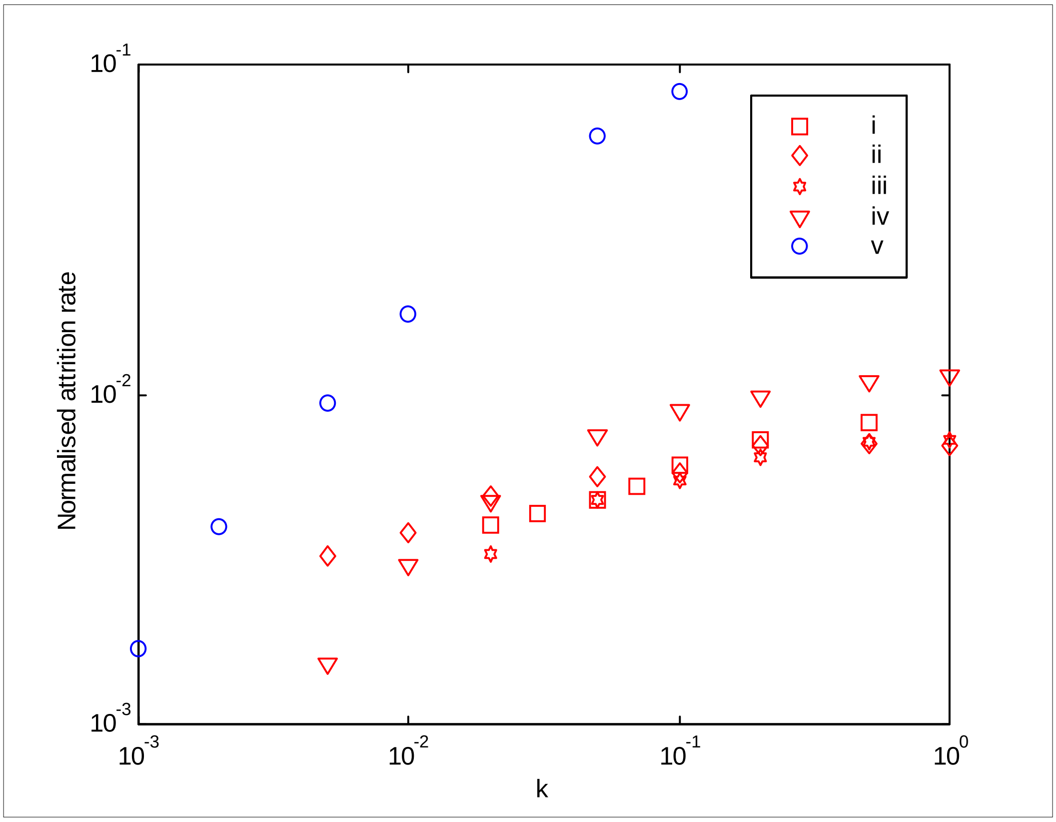

Figure 5: Mean attrition rate as a function of kill probability.

Figure 1) and case v is a stochastic Lanchester model. Power laws (hence q values) may be found by fitting straight lines on the log-log plot of this data. Note that above a certain value for k, the power law breaks down, with the slope tending to zero. This represents the case where the kill probability is sufficiently high that Blue is being killed almost immediately after encountering Red, so that increasing k has little further effect.

For cases i to iii, q was approximately 0.25, for case iv, $q \approx 0.7$, and for case v, $q \approx 0.9$. The relationship between the degree of concentration of Red firepower (characterised by D) and the power-law describing the dependence of attrition rate on k (characterised by q), conditional on reaching 25% casualties, is shown in Table 1. The data suggests an approximate relationship:

$$q = D/2 \tag{4.1}$$

Although the power-law exponents found for k from the data in Figure 5 do not match exactly D / 2, it must be recognised that the value of D is difficult to estimate and undoubtedly changes as both sides become depleted. Note that:

$$\Delta t^D \sim d^D \equiv \left(d^2\right)^{D/2} \equiv A^{D/2}$$

where A is an area mapped out by an automaton in a time Δ t. D /2 is then the fractal dimension obtained using the area of the box as the measure, rather than box width.

We found at least one other example[18] where equation 4.1 appears to hold. In that case, the Blue force was a reconnaissance group (four automata) attempting to find its way through a much larger (16 automata) dispersed Red force. For that case, $D = 0.79$ and $q = 0.36$. However, for the sake of brevity, further details of that study are not given here.

Case (requirement is to reach 25% casualty level).	D	q
Dispersed	0.7	0.28
Linear	1.7	0.86
Stochastic Lanchester	2	0.9
Recce	0.79	0.36

Table 1: Dependence of q on D.

5. Effect on distribution of outcomes

If the relationship in equation 4.1 holds, then we can interpret D as being a measure of the degree to which Red concentrates its forces for this ensemble. Note that whereas in cases ii and iii, Red's personality traits cause its forces to concentrate in this way, for case i the concentration is due to luck: that is, Red causes this level of casualties only when Blue is unfortunate enough to run into a concentration of Red. For all three cases, the important point is that for the ensembles considered, the effect on Blue attrition of lowering Red's k value is mitigated by the enhancement in firepower due to the concentration of Red forces. That is to say, Red does better for this ensemble than a Lanchester model would suggest (note from Figure 5 that for a Lanchesterian model, the relationship between attrition and k is near linear). However, this firepower concentration must be balanced against the gaps in the battlefield it creates, since clusters of Red automata must stay together for increasingly long periods to kill their target, leaving other areas unprotected. Thus an increasingly large ensemble of runs fail to reach 25 per cent casualties. Consequently, the nature of the combat is somewhat "bimodal", especially for low k values.

This bimodal nature is particularly evident when examining the differences between cases i, ii and iii. For case i and case iii, the model rarely reached 25 per cent Blue casualties (mean of 15 per cent) once Red k was set at less than 0.02. However, for case ii the same mean was obtained when k had a value of only 0.003, as is evident from Figure 4. The reason for this superior performance is Red's propensity to cluster for that case. However, allowing Blue to react to Red, as in case iii, improves Blue's performance. This is because the leading Blue elements concentrate firepower on Red automata in their path, creating gaps which the following elements may exploit. Whether or not Blue is able to penetrate the Red force depends on who is best able to concentrate their firepower in a given run. This has a significant effect on the distribution of outcomes. The distribution for case iii in Figure 4 suggests a bimodal nature, and has a much longer "tail" (i.e. more extreme results) than the other two cases.

In summary, it appears that if Red behaviour is set so that it deliberately clusters around Blue, it substantially improves its chances of killing a given fraction of Blue. Blue may counter this by attacking nearby Red, but Red will still sometimes get the better of it, inflicting much heavier casualties than is typical.

6. Discussion and conclusions

Using a cellular automaton model, some significant differences between this and Lanchester models have been highlighted in regard to modelling manoeuvre warfare.

It was found that by introducing automata "personalities" to the entities being modelled, a dispersed force would adapt a fractal distribution, reflecting the degree to which firepower was concentrated. Such a distribution is neither uniformly dispersed, nor tightly concentrated in space.

The Red automata in the cases examined were able to concentrate their firepower to improve the rate of attrition of the Blue force in circumstances where they lacked the firepower to destroy Blue on contact. At the same time, if Red needed to concentrate its forces for long periods of time to achieve this, Blue would have the opportunity to exploit the gaps created in the battlefield.

Selecting just those runs for which Red was able to cause some predetermined level of Blue casualties, the attrition rate for this ensemble obeyed a power-law dependence on kill probability. The power law exponent was apparently related to the fractal dimension of the distribution of forces. The value of the power-law exponent was typically around 1/3. By contrast, a stochastic Lanchester model produced a power-law exponent of nearly 1. Thus this ensemble does much better as the value of k is decreased than would be expected from a Lanchester model. What is so interesting about this is that the superior performance of this ensemble is not entirely due to luck (as it would be for a stochastic Lanchester model), but due to the way the automata have distributed themselves for each run in this ensemble.

Note that the linear case (case I in Figure 1) with attrition calculated in the same way as for case II behaved in a way which would be reasonably well approximated by the Lanchester equation for a range of conditions, showing that automata (or more generally, agent) models should not automatically be expected to produce different results from the traditional methods.

The apparent connection between the fractal dimension of the distribution of the forces, and the power-law dependence of the attrition rate on kill probability is an interesting extension on the work of earlier researchers into automaton models, which established that automatons tend to evolve into fractal distributions. Here D may be thought of as the degree to which Red is able to concentrate its firepower. This must be balanced by two things: firstly, the degree to which Red needs to remain dispersed to prevent itself from becoming an easy target for a long-range strike; secondly, the degree to which the need for Red to concentrate its firepower is balanced by the gaps this creates for Blue. While these two insights might be "intuitively" obvious, they do not emerge from either the stochastic or analytical variants of the Lanchester equation, without subjective involvement by the modeller.

The key result is that either Red concentrates its forces in such a way that it is able to destroy Blue at a faster rate than expected from a Lanchester model, or Blue exploits the gaps that result from that concentration. This has particularly interesting consequences for risk assessment. In particular, an apparently badly armed Red force can cause considerable damage to a superiorly armed Blue force if it is able to appropriately concentrate its forces. This is a particularly important consideration for peacekeeping activities in the Third World, where a poorly armed or trained militia which may not seem like much of a threat can in a coordinated attack cause a disaster for a well-equipped modern force. On the other hand, if the Blue force is able to react to such a concentration of force, it will often cause firepower "holes" in the Red distribution which it can exploit. Doing so will allow it to perform much better in a typical run. However, the distribution of outcomes becomes more complicated. Giving both sides the ability to react to each other's attempts to concentrate firepower leads to a much wider range of outcomes for the model, and may be somewhat bimodal, with Blue either being caught by a Red firepower trap, or being able to create firepower gaps which it can exploit. In short, treating one side (or both) as a "lifeless mass" fails to produce the full range of possible outcomes.

Future work should investigate the possibility of determining appropriate k values (hence armour to firepower ratios) for Blue and Red success. One possibility is that a certain level of Blue armour will usually force Red to concentrate its forces in an attempt to make use of its superior numbers, leaving gaps for Blue to exploit. If Red k is low enough, Blue will be able to exploit these gaps the majority of times. An interesting side question which arises is the extent to which armour can be traded off against mobility, since extra mobility effectively lowers Red k by allowing Blue to spend less time in Red's weapons range. Certainly, it is evident from the results presented that emphasis should be placed on aspects other than firepower and numbers, which are the prime parameters for the Lanchester model. An understanding of how firepower may be concentrated and manoeuvre exploited has the potential for much greater pay-offs than concentrating solely on these Lanchester parameters.

References

1. Lanchester FW (1914). Aircraft in Warfare: The Dawn of the Fourth Arm — No. V, the Principle of Concentration. *Engineering* **98:** 422-423.

2. Bellany I (1999). Modelling War. *Journal of Peace Research* **36**(6): 729-739.

3. Fricker RD (1998). Attrition Models of the Ardennes Campaign. *Naval Research Logistics* **45**(1): 1-21.

4. Gaver DP and Jacobs PA (1997). Attrition Modeling in the Presence of Decoys: An Operations-other-than-War Motivation. *Naval Research Logistics* **44**(5): 507-514.

5. Kress M and Talmor I (1999). A New Look at the 3:1 Rule of Combat through Markov Stochastic Lanchester Models. *Journal of the Operational Research Society* **50**(7): 733-744.

6. Lauren MK (1999). *Characterising the difference between Complex Adaptive and Conventional Combat Models*. DOTSE Report 169, NR 1335, (Defence Operational Technology Support Establishment, Auckland, New Zealand).

7. Rhodes JE (1998). in *Maneuver Warfare Science 1998*. Edited by F. G. Hoffman and G. E. Horne. (Published by the US Marine Corp. Combat Development Command).

8. Watson JF (1999). Living on the Edge. *The British Army Review* **123:** 19-29.

9. Scales RH (1999). Adaptive Enemies: Achieving Victory by Avoiding Defeat. *Joint Forces Quarterly* **23**, Fall.

10. McNaught KR (1999). The Effects of Splitting Exponential Stochastic Lanchester Battles. *Journal of the Operational Research Society* **50**(3): 244-254.

11. Hunt CW (1998). Uncertainty Factor Drives New Approach to Building Simulations. *Signal* **53** No 2: 75-77.

12. Ilachinski A (1997). Irreducible Semi-Autonomous Adaptive Combat (ISAAC): An Artificial-Life Approach to Land Warfare (U). CRM 97-61.10 (Centre for Naval Analyses).

13. Mandelbrot BB (1983). *The Fractal Geometry of Nature.* W. H. Freeman and Company.

14. Bak P, Chen K, and Creutz M (1989). Self-Organized Criticality and the 'Game of Life'. *Nature* **342:** 780-782.

15. Bak P, Chen K, and Tang C (1990). A Forest-Fire Model and Some Thoughts on Turbulence. *Physics Letters* **14**7: 297-300.`

16. Berlekamp ER, Conway JH, and Guy RK (1982). *Winning Ways for Mathematical Plays*, **2** (Academic).

17. Gardner M (1970). *Scientific American* **223** (4), 120-124; (5), 118; (6), 114.

18. Lauren, MK, and Baigent, DL, *Exploring the Recce-Counter Recce Scenario with ISAAC*, DOTSE Report, in production.

APPENDIX: ISAAC parameters.

Linear	Personalities: weighting towards ... (threshold range = 2)			
	Alive/Injured Red	Alive/Injured Blue	Red flag	Blue flag
Red	10	40	0	1
Blue	40	10	1	0
	Meta personalities			
	Cluster	Advance	combat	
Red = Blue	10	5	3	
	Attributes			
	Fire range	Sensor range	Movement rate	*k* (baseline)
Red = Blue	3	5	2	0.05
	Initial distribution (battlefield size 120x120)			
	Centre position		Size	
Red	30,30		15,15	
Blue	90,90		15,15	
Flag	60		60	

Table 2: Parameters for case I.

	Personalities: weighting towards ... (threshold range = 5)			
	Alive/Injured Red	Alive/Injured Blue	Red flag	Blue flag
Red i)	0	0	0	0
ii)	10	40	0	0
iii)	10	40	0	0
Blue i)	0	10	10	0
ii)	0	10	10	0
iii)	40	10	10	0
	Meta personalities			
	Cluster	Advance	combat	
Red = Blue	3	1	3	
	Attributes			
	Fire range	Sensor range	Movement rate	*k* (baseline)
Red = Blue	3	5	1	0.05
	Initial distribution (battlefield size 120x120)			
	Centre position	Size		
Red	60,60	30,30		
Blue	60,105	15,15		
Flag	60	1		

Table 3: ISAAC parameters for Case II.